

AD 651868

SEMI-MARKOV PROCESSES: A PRIMER

Bennett Fox

April 1967

DDC
RECEIVED
MAY 22 1967
B

STATEMENT NO. 1

Distribution of This Document is Unlimited

P-1577

ARCHIVE COPY

200409 10020

SEMI-MARKOV PROCESSES: A PRIMER

Bennett Fox^{*}

The RAND Corporation, Santa Monica, California

Marrying renewal processes and Markov chains yields semi-Markov processes, and the former are special cases of the latter. In this expository paper, some of the main properties of the union are outlined.

* Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

Invited paper prepared for submission to Management Science.

PER 07 96000

I. INTRODUCTION

Our description of semi-Markov processes (SMP's) is heuristic and unconventional. We postpone a formal definition of a distinguished state and a discussion of regularity conditions to following sections. Many recondite matters, especially those that arise in the infinite state case, are omitted.

A key concept, already familiar to readers acquainted with queueing theory, is that of an imbedded Markov chain connecting the regeneration points of a stochastic process. Starting from a regeneration point, the future is stochastically independent of the past. All distinguished states correspond to regeneration points; thus, they have the Markov property. In the imbedded Markov chain, the original time scale of the transitions between the distinguished states is replaced by a discrete time version where all transitions take unit time. By studying the corresponding semi-Markov process, defined precisely in Sec. 3, we recover the probabilistic behavior in the original time scale.

I^+ is defined to be the set of distinguished states, assumed countable. The state transitions form a Markov chain with transition probabilities (p_{ij}) , where direct transitions from a state to itself (e.g., $p_{ii} > 0$) are allowed. Given that an $i \rightarrow j$ transition is about to occur, the duration of the transition has distribution F_{ij} . We define*

$$Q_{ij}(t) = p_{ij}F_{ij}(t)$$

$$H_i(t) = \sum_j Q_{ij}(t).$$

*Unless otherwise stated, all summations will be over I^+ and all functions vanish for negative arguments.

Thus H_i is the unconditional distribution of the time elapsed starting from state i until the next state is entered (possibly i itself). Note that

- (i) a one-state SMP is a renewal process;
- (ii) an SMP with F_{ij} degenerate at one for all i, j is a Markov chain;
- (iii) an SMP with all F_{ij} exponential is a continuous time countable-state Markov process.

Many problems in management science and operations research can be modeled as SMP's: for example, queueing, inventory, and maintenance problems. For details, see, e.g., Pyke [27], Barlow and Proschan [1], Fabens [11], and Çinlar [4]. For proofs, citations of earlier papers, and additional topics in SMP's, the reader should consult the reference list. Another expository paper is Janssen [19].

II. DISTINGUISHED STATES

In this section we examine precisely what is meant by a distinguished state. Since discussing this topic in an offhand manner could result in confusion, the subject is treated in some detail.

For each sample path of a stochastic process X , there is a correspondence between $[t: t \geq 0]$ and a set of states S . If every state in S is required to have the Markov property, then in general S will be uncountable since a history of the process, or at least the relevant portion of it, must be part of the state definition. However, all that we require of S is that it have an appropriate countable subset I^+ of (distinguished) states having the Markov property. Thus, for a state to be a candidate for I^+ , it must correspond to a regeneration point, but we do not require that all states corresponding to regeneration points belong to I^+ .

To fix these ideas concretely, we illustrate them with examples from the M/G/1 queue (Poisson arrivals, general service time distribution, single channel). For many purposes, a convenient set of distinguished states is $\{0, 1, 2, \dots, i, \dots\}$ where state i signifies i customers in the system and a service has just been completed. Note that not all regeneration points are included in this set, since any time the system is idle (empty), it is at a regeneration point. Our choice of distinguished states conforms to our requirements because the time to the next arrival is stochastically independent of the time elapsed since the last arrival. In general, arrival epochs, except those corresponding to the start of a busy period, are not regeneration points. Thus the state " i customers in the system and

a customer has just arrived" cannot be a distinguished state, unless the service times are exponential. Through the use of so-called supplementary variables, we can define every state in the original process so that it is Markovian. Each state is then a couple of the form (i, u) , which denotes i customers in the system and the customer being processed has been in service for time u . Sometimes supplementary variable techniques are useful as an alternative or adjunct to SMP techniques; see, e.g., Cox and Miller [5].

Returning to the general discussion, we require the distinguished states to be defined such that non-zero holding times in a distinguished state are forbidden but instantaneous transitions among the distinguished states are allowed. This is a departure from the setup of Pyke [26], although the two formulations are essentially equivalent. Our definition of distinguished state is natural in a dynamic programming framework (see, e.g., Denardo and Mitten [9]) and permits a graphic representation of SMP's in terms of networks with branch nodes (distinguished states) and stochastic arc lengths; see [13]. For example, traversing an arc could correspond to a customer completing service.

To remove ambiguity in case of instantaneous transitions,* we define $X^+(t) = X(t^+)$; thus, X^+ is right continuous and last distinguished state of X^+ entered in $[0, t]$, say, is well defined. Since we have prohibited non-zero holding times in distinguished states, we cannot allow a distinguished state to correspond to a nondegenerate interval of regeneration points (e.g., an idle period in an M/G/1 queue). Thus,

* Sometimes it is convenient to permit instantaneous transitions; see, e.g., Denardo [7].

to exclude an infinite sequence of instantaneous transitions from a state to itself, we require that the distinguished states be defined such that, for all nondegenerate intervals (a, b) , $i \in I^+ \Rightarrow P\{X(t) = i, \forall t \in (a, b)\} = 0$. For example, in the M/G/1 queue it does not suffice to define the distinguished state 0 as 0 customers in the system. The condition that a service has just been completed must be added.

Throughout the sequel, "state" refers only to "distinguished state," necessitating definitions slightly different from conventional usage.

III. REGULAR SMP'S

Letting $N_j(t)$ denote the number of times state j is entered in the half-open interval $(0, t]$, we obtain the Markov Renewal Process (MRP) $\underline{N}(t) = (N_0(t), N_1(t), \dots, N_n(t))$, where $I^+ = \{0, 1, 2, \dots, n\}$, possibly with $n = \infty$. Let $Z(t)$ be the last distinguished state of X^+ entered in $[0, t]$. In the literature the Z process is called an SMP. However, the MRP and the SMP are different aspects of the same underlying stochastic process; therefore, by slight abuse of language, we shall refer to the underlying process itself as an SMP or a MRP, using the terms interchangeably.

A MRP is regular if with probability one (w.p.1) each state is entered only a finite number of times in any finite time span--i.e., if $P[N_i(t) < \infty] = 1, \forall i \in I^+$ and $t \geq 0$. A MRP is strongly regular if w.p.1 the total number of state transitions is finite in any finite time span--i.e., if $P[\sum N_i(t) < \infty] = 1, \forall t \geq 0$. Clearly strong regularity implies regularity and, if $n < \infty$, the terms are equivalent. In applications, it will ordinarily be obvious that strong regularity holds. If $n < \infty$, it suffices that $\underline{H} = (H_0, \dots, H_n)$ have at least component nondegenerate at zero for every ergodic sub-chain of the imbedded Markov chain. In the denumerable state case ($n = \infty$), see Pyke [26] and Pyke and Schaufele [27] for conditions that imply strong regularity. In the sequel, we assume that strong regularity holds.

IV. FIRST PASSAGE AND COUNTING DISTRIBUTIONS

Let

$$P_{ij}(t) = P[Z(t) = j | Z(0) = i]$$

$$G_{ij}(t) = P[N_j(t) > 0 | Z(0) = i]$$

(first passage time distribution)

$$M_{ij}(t) = E[N_j(t) | Z(0) = i]$$

(mean entry counting function).

Defining the convolution

$$(A * B)(t) = \int_0^t A(t-x) dB(x),$$

we have by straightforward renewal-theoretic arguments:

$$P_{ij} = (1 - H_i)\delta_{ij} + \sum_k Q_{ik} * P_{kj} = (1 - H_i)\delta_{ij} + P_{jj} * G_{ij}$$

$$G_{ij} = Q_{ij} + \sum_{k \neq j} Q_{ik} * G_{kj}$$

$$M_{ij} = G_{ij} + G_{ij} * M_{jj} = Q_{ij} + \sum_k Q_{ik} * M_{kj}.$$

In general, these relations cannot be solved analytically. In the finite state case, numerical solutions can be obtained by numerical inversion of the corresponding Laplace transforms (see, e.g., [2]).

[18], [24], and [30]). For each value of s , only one matrix inversion in the transform domain is required--that of $I - q(s)$, where $s > 0$ and

$$q(s) = \left(\int_{0^-}^{\infty} e^{-st} dQ_{ij}(t) \right).$$

In obvious notation, having found $[I - q(s)]^{-1}$, either analytically as a function of s or, for suitably spaced values of s , numerically, one successively computes

$$m(s) = [I - q(s)]^{-1} q(s) = [I - q(s)]^{-1} - I$$

$$g_{ij}(s) = m_{ij}(s) / [1 + m_{jj}(s)]$$

$$p_{ij}(s) = p_{jj}(s) g_{ij}(s), \quad i \neq j$$

$$p_{jj}(s) = \frac{1 - h_j(s)}{1 - g_{jj}(s)},$$

and then inverts the transforms. Although this procedure is not trivial, it often compares favorably with the alternative simulation approach for getting the transient behavior in the time domain. By usual limit theorems for Laplace transforms (Widder [31], Feller [12], see also Jewell [20]), the behavior in the time domain for large (small) t corresponds to behavior in the transform domain for small (large) s .

Conditioning on the event that no state in a subset B of I^+ is entered in $(0, t]$ may be of interest. For example,

$${}_B P_{ij}(t) = P[Z(t) = j | Z(0) = i, N_k(t) = 0, \forall k \in B]$$

$${}_B G_{ij}(t) = P[N_j(t) > 0 | Z(0) = i, N_k(t) = 0, \forall k \in B]$$

$${}_B M_{ij}(t) = E[N_j(t) | Z(0) = i, N_k(t) = 0, \forall k \in B]$$

can be calculated from the formulas already given by (temporarily) making the states in B absorbing.

The first and second moments of G_{ij} , denoted respectively by μ_{ij} and $\mu_{ij}^{(2)}$, are given by

$$\mu_{ij} = \sum_{k \neq j} p_{ik} \mu_{kj} + v_i$$

$$\mu_{ij}^{(2)} = \sum_{k \neq j} p_{ik} [\mu_{kj}^{(2)} + 2v_{ik} \mu_{kj}] + v_i^{(2)},$$

where v_{ij} is the mean of F_{ij} , $v_i^{(1)} = v_i$, and

$$v_i^{(2)} = \int_0^\infty t^2 dH_i(t).$$

We assume that $v_i^{(2)} < \infty$, $\forall i \in I^+$. If the imbedded Markov chain is finite and ergodic, these equations have a unique finite solution* and, with π_j the stationary probability that the last state entered

* See appendices 1 and 2 of Fox [13] for an efficient way to solve these equations. (An expression for the "bias terms" in Markov renewal programming involves the first passage time moments, which are of intrinsic interest, but recently Jewell [22] has derived a remarkably simple alternative expression, obviating the need to calculate $\{\mu_{ij}\}$ and $\{\mu_{ij}^{(2)}\}$ to evaluate the bias terms).

is j if all F_{ij} were degenerate at one, * multiplying μ_{ij} and $\mu_{ij}^{(2)}$ by π_i and summing yields

$$\mu_{jj} = (1/\pi_j) \sum_k \pi_k v_k$$

$$\mu_{jj}^{(2)} = (1/\pi_j) \left[\sum_k \pi_k v_k^{(2)} + 2 \sum_{k \neq j} \pi_i p_{ik} v_{ik} \mu_{kj} \right].$$

For finite state SMP's, the probability that state j is ultimately reached starting from i is **

$$G_{ij}(\infty) = \begin{cases} 1, & \text{if } i, j \in E_k \\ 0, & \text{if } i \in E_k, j \in E_l, k \neq l \\ [(I - A)^{-1} \theta]_i, & \text{if } i \in T, j \in E_k \end{cases}$$

where A is the submatrix of P corresponding to the set T of transient states, E_1, \dots, E_m are the recurrent subchains of P , and

$$\theta_i = \sum_{l \in E_k} p_{il}, \quad i \in T.$$

The mean time to leave T starting from i is

$$\xi_i = [(I - A)^{-1} v^t]_i, \quad i \in T,$$

where v^t is the vector of v_j 's, $j \in T$. A double generating function for the distribution of $N_j(t)$ is

* In other words, π is the stationary measure for the imbedded chain, but not (in general) for the SMP itself.

** The case $i, j \in T$, of less interest, is not considered.

$$Y_z = \underline{1} - (1 - z)m [z I + (1 - z)D]^{-1},$$

where $\underline{1}$ is a matrix of 1's,

$$D = \left((1 - q_{ij}) \delta_{ij} \right)^{-1} = \left(\frac{\delta_{ij}}{1 - q_{ij}} \right)$$

and

$$Y_z = (\phi_{ij}(z; \bullet))$$

$$\phi_{ij}(z; s) = \int_0^\infty e^{-st} d_t w_{ij}(z; t)$$

$$w_{ij}(z; t) = \sum_{k=0}^{\infty} z^k v_{ij}(k; t)$$

$$v_{ij}(k; t) = P[N_j(t) = k | Z(0) = i].$$

Thus, in principle, the probabilities and moments can be obtained in the usual way. The Laplace transform m of the first moment $(M_{ij}(\bullet))$ was already given. We remark that, if $Z(0) = i$ and i belongs to the same ergodic subchain as j , $t^{-1}N_j(t) \rightarrow 1/\mu_{jj}$ ($1/\infty \equiv 0$) w.p.1, a strong law that follows immediately from renewal theory. See Pyke and Schaufele [28] for further general moment computations, weak and strong laws of large numbers, and central limit theorems.

V. STATE CLASSIFICATION

In classifying the states of a SMP transient, null recurrent, or positive recurrent, we must distinguish between a state's classification in the imbedded Markov chain and in the SMP itself. For I^+ finite and $v_i < \infty$, $\forall i \in I^+$, the distinction disappears and a state j is either transient or positive recurrent ($G_{jj}(\infty) = 1$ and $\mu_{jj} < \infty$). In large-scale applications, the ergodic subchain-transient set breakdown may not be obvious and recourse may be necessary to an algorithmic classification scheme such as that of Fox and Landi [16].

For I^+ infinite, a state j is transient (recurrent--i.e., $G_{jj}(\infty) = 1$) in the SMP $\Leftrightarrow j$ is transient (recurrent) in the imbedded Markov chain. State l is positive recurrent in the imbedded Markov chain (contained in ergodic subchain E_k) and, for some constant c , $v_{ij} \leq c < \infty$, $\forall i, j \in E_k$, $\Rightarrow l$ is positive recurrent in the SMP. A SMP is positive recurrent if all the states in I^+ are positive recurrent in the SMP.

VI. STATIONARY PROBABILITIES

It is important to distinguish between the stationary probabilities $\{\pi_i\}$ with respect to the imbedded Markov chain and the stationary probabilities $\{\rho_i\}$ with respect to the SMP.* Thus ρ_i is the steady state probability that the last distinguished state entered is i . Hence the $\{\rho_i\}$ are of direct interest in applications, while the $\{\pi_i\}$ are computed only as an intermediate step. We consider first the case I^+ finite and $v_i < \infty$, $\forall i \in I^+$.

$$\rho_j = \begin{cases} v_j/\mu_{jj}, & j \in E_k, Z(0) \in E_k \\ G_{ij}(\infty)v_j/\mu_{jj}, & j \in E_k, Z(0) = i \in T \\ 0, & j \in T \\ 0, & j \in E_k, Z(0) \in E_l, k \neq l \end{cases}$$

where $G_{ij}(\infty)$ was computed already and

$$v_j/\mu_{jj} = \frac{\pi_j v_j}{\sum_{i \in E_k} \pi_i v_i}$$

with $\{\pi_i\}$ here being the stationary probabilities for the imbedded Markov chain given that $Z(0) \in E_k$.

In the remainder of this section we assume that the imbedded Markov chain is irreducible and that the SMP is positive recurrent, where I^+ may be finite or infinite. We also assume that the mean

*In general, the stationary probabilities must be interpreted as Cesaro limits. If the process is aperiodic, these reduce to ordinary limits.

transition times are uniformly bounded away from zero; i.e., $0 < e \leq v_{ij} < \infty$. With these assumptions, Fabens [10] shows that

$$\rho_j = \frac{\pi_j v_j}{\sum_i \pi_i v_i}$$

in agreement with result given above for the I^+ finite case. Define $\sigma(x)$ as the time of the last transition before or at x and $\tau(x)$ as the time of the next transition after x . The random variables

$$\gamma(x) = \tau(x) - x \quad (\text{excess r.v.})$$

$$\delta(x) = x - \sigma(x) \quad (\text{shortage r.v.})$$

are of interest. Adding to the previous assumptions the hypotheses that the mean recurrence times $\{\mu_{ii}\}$ are finite and that $Z(e)$ is aperiodic, Fabens shows that

$$\lim_{t \rightarrow \infty} P\{\delta(t) \leq x | Z(t) = i\} = \lim_{t \rightarrow \infty} P\{\gamma(t) \leq x | Z(t) = i\}$$

$$= \frac{1}{v_i} \int_0^x [1 - H_i(u)] du.$$

This generalizes the well-known result from renewal theory for the one state case, obtained there as a corollary to the key renewal theorem; see, e.g., Barlow and Proschan [1].

The general question of existence and uniqueness of stationary measures is dealt with in Pyke and Schaufele [29]. Oneong [3] gives conditions under which convergence to the steady state is geometric.

VII. EXAMPLE: THE M/G/1 QUEUE

To illustrate the notion of stationary probabilities for a SMP, we consider the M/G/1 queue. Let

π_i = the stationary probability that i customers are in the system just after a random service completion epoch

ρ_i = the stationary probability that i customers are in the system just after the service completion epoch preceding a random point in time

p_i = the stationary probability that i customers are in the system at a random point in time.

We assume that the traffic intensity λb is less than one, where λ is the arrival rate and b the mean service time, assumed positive.

Although it is easily shown that

$$p_i = \begin{cases} \lambda b \pi_i, & i \geq 1 \\ (1 + \lambda b) \pi_0 = 1 - (\lambda b)^2, & i = 0, \end{cases}$$

it turns out that $p_i = \pi_i$, $\forall i$, a remarkable result originally due to Khintchine [23] and derived in a more elementary manner by Fox and Miller [17] using SMP theory. A similar result holds for the G/M/1 queue, but in bulk queues (Fabens [11]), for example, the stationary measures for the imbedded Markov chain and the original queueing process are different.

Readers familiar with queueing theory may prefer to skip to the last paragraph of this section. In between, the standard manipulations yielding $G_\pi(z)$, the generating function of the $\{\pi_i\}$, are performed.

Recalling that state n means that there are n people in the system and a service has just been completed, we obtain the well known

transition matrix for the imbedded Markov chain:

$P =$

	0	1	2	3	4	.	.	.
0	k_0	k_1	k_2	k_3	k_4	.	.	.
1	k_0	k_1	k_2	k_3	k_4	.	.	.
2	0	k_0	k_1	k_2	k_3	.	.	.
3	0	0	k_0	k_1	k_2	.	.	.
4	0	0	0	k_0	k_1	.	.	.
.
.
.

where

$$k_n = \int_0^{\infty} \frac{e^{-\lambda t}}{n!} (\lambda t)^n dB(t)$$

and B is the service distribution. By the usual straightforward manipulations, we find that the generating function of the $\{k_n\}$ is

$$G_k(z) = \sum_i k_i z^i = \beta(\lambda(1-z)),$$

where β is the Laplace-Stieltjes transform of B, i.e.,

$$\beta(s) = \int_0^{\infty} e^{-st} dB(t).$$

To obtain the stationary vector π for the chain, we multiply the i -th relation determined by $\pi P = \pi$ by z^i and sum, define the generating function

$$G_{\pi}(z) = \sum_i \pi_i z^i,$$

and obtain from the special form of P for this chain by an easy calculation the standard result

$$G_{\pi}(z) = \frac{\pi_0(1-z)G_k(z)}{G_k(z) - z}.$$

Using the fact that $\lim_{z \rightarrow 1^-} G_{\pi}(z) = 1$ (i.e., the probabilities sum to 1) and applying L'Hospital's rule,

$$\pi_0 = 1 - \lambda b.$$

Summarizing our results so far,

$$G_{\pi}(z) = \frac{\pi_0(1-z)\beta(\lambda(1-z))}{\beta(\lambda(1-z)) - z}.$$

Thus the mean number in the system averaged over service completion epochs is, with σ^2 the variance of the service times,

$$\lim_{z \rightarrow 1^-} G_{\pi}(z) = \pi_0 \lim_{z \rightarrow 1^-} \left[\frac{G_k''(z) - 2G_k'(z)(G_k'(z) - 1)}{2(G_k'(z) - 1)^2} \right] = \frac{(\lambda b)^2 + \sigma^2}{1 - \lambda b} + \lambda p,$$

and by the fact that $\pi_i = p_i$, $\forall i$, is also the mean number in the system at random point in time (in the steady state). Higher moments

and probabilities can be obtained from the generating function by appropriate differentiations, which, however, become quite tedious.

Having found $G_{\pi}(z)$, the stationary waiting distribution for the first come, first served (FIFO) discipline can easily be found. If an arrival finds the system empty, the conditional wait in queue is 0. Otherwise it is governed by the remaining processing time of the customer in service, the excess random variable, plus the service times for the customers (if any) already in queue. Noting that $G_p(z) = G_{\pi}(z)$, the interested reader can readily derive the Pollaczek-Khintchine formula for the stationary wait in queue. (The Laplace-Stieltjes transform of the stationary queueing delay distribution is found to be $\pi_0 + \frac{1 - \beta(s)}{s\beta(s)} [G_p(\beta(s)) - \pi_0]$.)^{*} See Feller [12], p. 392, for an alternate elegant derivation that bypasses the calculation of $G_p(z)$. A third derivation follows from the fact that the number of customers in the system just after a departure is the number of arrivals during his total wait (queueing time plus service time); the resulting equation is solved by taking generating functions yielding the standard form of the Laplace-Stieltjes transform of the stationary queueing delay distribution $s\pi_0/(s - \lambda(1 - \beta(s)))$, a version of the Pollaczek-Khintchine formula. For a fourth derivation, where the (superfluous) assumption of a continuous failure distribution is tacitly made, see Cox and Miller [5], pp. 241-242.

^{*} Comparison with the standard form of the transform yields an interesting and surprising identity. The derivation depends on the fact that, since the arrival process is Poisson, an arrival plays the role of a random observer.

VIII. ASYMPTOTIC FORM OF M_{ij}

From a previous section, we know that for finite state SMP's

$$\begin{aligned} m_{ij}(s) &= g_{ij}(s)[1 + m_{jj}(s)] \\ &= g_{ij}(s) \left[1 + \frac{g_{jj}(s)}{1 - g_{jj}(s)} \right]. \end{aligned}$$

Formally expanding e^{-sx} in a Taylor series and integrating termwise yields for i, j in the same ergodic subchain

$$m_{ij}(s) = \frac{1}{s\mu_{jj}} + \frac{\mu_{jj}^{(2)}}{2\mu_{jj}^2} - \frac{\mu_{ij}}{\mu_{jj}} + o(1),$$

whence by a Tauberian argument

$$M_{ij}(t) - \frac{t}{\mu_{jj}} \xrightarrow[\text{(Cesàro)}]{(2)} \frac{\mu_{jj}^{(2)}}{2\mu_{jj}^2} - \frac{\mu_{ij}}{\mu_{jj}},$$

a result that can be obtained by analogy with renewal theory for delayed recurrent events, where the time to the first "renewal" has distribution G_{ij} and the spacing between subsequent renewals has distribution G_{jj} . If the SMP is aperiodic, the Cesaro limit reduces to an ordinary limit. It can be shown that the formal manipulation used to obtain the asymptotic expansion of $m_{ij}(s)$ is justified provided that $v_i^{(2)} < \infty$, $\forall i \in I^+$.

IX. FINITE SMP'S WITH COSTS

Often in applications, costs are associated with the transitions. Measuring time from the start of an $i \rightarrow j$ transition, let $C_{ij}(x|t)$ be the cost incurred up to time x given that the transition length is t . The expected discounted cost for a transition starting from state i is then

$$\gamma_i(\alpha) = \sum_j p_{ij} \int_0^\infty dF_{ij}(t) \int_0^t e^{-\alpha x} d_x C_{ij}(x|t),$$

where a cost incurred at time x is discounted by the factor $e^{-\alpha x}$.

An elementary renewal type argument then shows that $v_i(\alpha)$, the total expected discounted cost over an infinite horizon starting from state i , satisfies

$$v(\alpha) = \gamma(\alpha) + q(\alpha)v(\alpha),$$

where $\alpha > 0$ and $v(\alpha)$ and $\gamma(\alpha)$ are the vectors with components $v_i(\alpha)$ and $\gamma_i(\alpha)$, respectively. Thus, assuming a finite number of states,

$$v(\alpha) = [I - q(\alpha)]^{-1} \gamma(\alpha).$$

Since $I - q(0)$ is singular and a direct asymptotic expansion is not obvious, it is convenient to make use of the relation between q and m given earlier to study the behavior of $v(\alpha)$ as $\alpha \rightarrow 0^+$. Following Jewell [21], we have

$$v(\alpha) = [I + m(\alpha)] \gamma(\alpha).$$

Making use of the expansion of $m(\alpha)$ given in the preceding section, we find that, if i is a recurrent state, $v_i(\alpha)$ has the form

$$v_i(\alpha) = l_i/\alpha + w_i + o(1)$$

and a straightforward argument in Fox [13] then shows that this form is valid for any state; i.e.,

$$v(\alpha) = l/\alpha + w + o(1) ,$$

where expressions for l , the loss rate vector, and w , the bias term vector, can be found in Jewell [21, 22] and Fox [13], where appropriate conditions are given to justify the expansion. Substituting this relation into $v(\alpha) = Y(\alpha) + q(\alpha)v(\alpha)$ and equating the coefficients of α^{-1} and the constant terms, respectively, yields*

$$Pl = l$$

$$Y + Pw = w + y$$

$$y_i = \sum_j p_{ij} v_{ij} l_j$$

$$= v_i l_i, \quad \text{if } i \text{ is recurrent.}$$

These expressions can be solved uniquely for l , but w is determined only up to an additive constant in each ergodic subchain; see, e.g.,

*This procedure can be justified by a simple contradiction argument. Note that $q_{ij}(\alpha) = p_{ij}(1 - \alpha v_{ij}) + o(\alpha)$ and that the loss rate for all states in an ergodic subchain is the same.

Denardo and Fox [8]. An interesting and intuitive result that follows easily from the above formulas is that the loss rate for each state in an ergodic subchain E_k is the same and equal to $\pi^{(k)} \gamma' / \pi^{(k)} \nu'$, where $\pi^{(k)}$ is the stationary vector for the corresponding submatrix and γ' and ν' are the restrictions γ and ν , respectively, to E_k . The loss rates for the transient states are obtained from the fact reflected in $Pf = f$ that the loss rate for a state is given by the appropriate convex combination and that $I - A$, where A is the submatrix corresponding to the transient states, is invertible.

Denoting the undiscounted loss up to time t by $L(t)$, we obtain from the asymptotic expansion of $v(\alpha)$ that

$$L(t) - \lambda t \xrightarrow[\text{(Cesàro)}]{} w.$$

X. MARKOV RENEWAL PROGRAMMING

The situation becomes more interesting when, at each state i , one has a set of options A_i and the choice at i simultaneously determines p_{ij} , F_{ij} , and C_{ij} for all $j \in I^+$. The goal is either to choose a policy that minimizes either the expected discounted loss or the loss rate. In the latter case, an appropriate secondary objective is to minimize (L, w) lexicographically, which is especially important when some policies can have transient states. With either criterion, an optimal policy can be found by linear programming when I^+ and $\sum_{i \in I^+} A_i$ are finite. For details, see, e.g., Jewell [21, 22], Fox [13], Denardo [6, 7], and Denardo and Fox [8], where references to the earlier (extensive) literature on the subject are given. Some papers (e.g., Derman [10]) treat the I^+ infinite case, but the author believes that for applications the general theory developed so far for that case is inadequate and that particular problems are best attacked on an ad hoc basis. For the case where I^+ is finite but the finiteness restriction on $\sum A_i$ is dropped, see Fox [14].

XI. ESTIMATION AND STATISTICAL INFERENCE

Moore and Pyke [25] develop estimators for the $\{p_{ij}\}$ and the $\{F_{ij}\}$ and their large sample distributions. For statistical inference in birth and death queueing models, see Wolff [32]. Both of the foregoing approaches are objectivist, i.e., non-Bayesian. When a large number of observations are at hand, the objectivist approach is unobjectionable and difficulties stemming from a possible lack of consensus of prior belief do not emerge. On the other hand, when the observations are few or nonexistent, as is common, a Bayesian approach incorporating prior beliefs and loss functions is essential.* Such an approach may be formal or may simply consist of a sensitivity analysis with the outcomes being given subjective weights. In the realm of decision making, policies should adapt to modified beliefs as more observations are taken.** This area remains largely unexplored and is ripe for investigation.

* This is, of course, a statement of the author's opinion. These matters are highly controversial.

** In this connection, [15] may be of interest.

REFERENCES

1. Barlow, R. E., and F. Proschan, Mathematical Theory of Reliability, John Wiley and Sons, New York, 1965.
2. Bellman, R., H. Kagiwada, R. Kalaba, and M. Prestrud, Invariant Imbedding and Time Dependent Transport Processes, American Elsavier, New York, 1964.
3. Cheong, C. K., "Geometric Convergence of Semi-Markov Transition Probabilities," Z. Wahrscheinlichkeitstheorie verw. Geb., Vol. 7 (1967), pp. 122-130.
4. Çinlar, E., "Decomposition of a Semi-Markov Process under a Markovian Rule," Austral. J. Statist., Vol. 8 (1966), pp. 163-179.
5. Cox, D. R., and H. D. Miller, The Theory of Stochastic Processes, John Wiley and Sons, New York, 1965.
6. Denardo, E. V., "Contraction Mappings in the Theory Underlying Dynamic Programming," to appear in SIAM Review.
7. -----, Separable Markovian Decision Problems, P-3448, The RAND Corporation, Santa Monica, California, September 1966.
8. -----, and B. L. Fox, Multichain Markov Renewal Programs, RM-5208-PR, The RAND Corporation, Santa Monica, California, February 1967.
9. -----, and L. G. Mitten, "Elements of Sequential Decision Processes," J. of Indust. Engin., Vol. 18 (1967), pp. 106-112.
10. Derman, C., "Denumerable State Markovian Decision Processes--Average Cost Criterion," Ann. Math. Stat., Vol. 37 (1966), pp. 1545-1553.
11. Fabens, A. J., "The Solution of Queueing and Inventory Models by Semi-Markov Processes," J. Roy. Statist. Soc. Series B, Vol. 23 (1961), pp. 113-127.
12. Feller, W., An Introduction to Probability Theory and Its Applications, Vol. 2. John Wiley and Sons, 1966.
13. Fox, B. L., "Markov Renewal Programming by Linear Fractional Programming," SIAM J. Appl. Math., Vol. 14, pp. 1418-1432.
14. -----, "Existence of Stationary Optimal Policies for Some Markov Renewal Programs," to appear in SIAM Review.
15. -----, "Adaptive Age Replacement," to appear in J. Math. Anal. and Appl.

16. Fox, B. L., and D. M. Landi, An Algorithm for Identifying the Ergodic Subchains and Transient States of a Stochastic Matrix, RM-5269-PR, The RAND Corporation, Santa Monica, California, March 1967.
17. -----, and L. W. Miller, Equivalent Stationary Measures for the M/G/1 Queue, P-3590, The RAND Corporation, Santa Monica, California, April 1967.
18. Gaver, D. P., "Observing Stochastic Processes and Approximate Transform Inversion," Operations Research, Vol. 14 (1966), pp. 444-459.
19. Janssen, J., "Processus de Renouvellements Markoviens et processus Semi-markoviens," Cahiers du C.E.R.O.: Part 1, Vol. 6 (1964), pp. 81-105; Part 2, Vol. 7 (1965), pp. 126-141.
20. Jewell, W. S., "The Properties of Recurrent Event Processes," Operations Research, Vol. 8 (1960), pp. 446-472.
21. -----, "Markov-Renewal Programming. I and II," Operations Research, Vol. 11 (1963), pp. 938-971.
22. -----, "A Note on Markov Renewal Programming," ORC 66-41, Operations Research Center, University of California, Berkeley, November 1966.
23. Khintchine, A. J., "Mathematisches über die Erwartung vo. einem öffentlichen Schalter," Mat. Sbornik, Vol. 39 (1932), pp. 73-84 (in Russian, German summary).
24. Miller, M. K., and W. T. Guy, "Numerical Inversion of the Laplace Transform by Use of Jacobi Polynomials," SIAM J. Numer. Analysis, Vol. 3 (1966), pp. 624-635.
25. Moore, E. H., and R. Pyke, "Estimation of the Transition Distributions of a Markov Process," D1-82-0371, Boeing Scientific Research Laboratories, Seattle, Washington, September 1964.
26. Pyke, R., "Markov Renewal Processes: Definitions and Preliminary Properties," Ann. Math. Stat., Vol. 32 (1961), pp. 1231-1242.
27. -----, "Markov Renewal Processes with Finitely Many States," Ann. Math. Stat., Vol. 32 (1961), pp. 1243-1259.
28. -----, and R. Schaufele, "Limit Theorems for Markov Renewal Processes," Ann. Math. Stat., Vol. 35 (1964), pp. 1746-1764.
29. -----, and -----, "The Existence and Uniqueness of Stationary Measures for Markov Renewal Processes," Ann. Math. Stat., Vol. 37 (1966), pp. 1439-1462.

30. Spinelli, R. A., "Numerical Inversion of a Laplace Transform," SIAM J. Numer. Analysis, Vol. 3 (1966), pp. 636-649.
31. Widder, D. V., The Laplace Transform, Princeton University Press, Princeton, 1946.
32. Wolff, R. W., "Problems of Statistical Inference for Birth and Death Queueing Models," Operations Research, Vol. 13 (1965), pp. 343-357.